

What is claimed is:-

1. A method of indexing content in an IP-based network comprising:-
 - (a) intercepting traffic flowing in the network,
 - (b) extracting content identity information and associated destination location information from the traffic flow,
 - (c) generating a mapping from a content item identified by the extracted identity information to at least one destination location identified by the associated destination location information, and
 - (d) storing the mapping in a content index database which is operable to provide an instance mapping containing a list of destination locations in response to an instance request containing a content identity.
2. A method according to claim 1, wherein the step of intercepting traffic is carried out by intercepting traffic flowing into a cache, and wherein the method further comprises advertising the content identities for which mappings are stored in the content index by sending advertising messages to a predetermined location in the network.
3. A method according to claim 2, wherein the method further comprises recording the time of traffic flows into the cache which are related to a particular content item and calculating the time period between a first flow of the content item into the cache and a subsequent flow of the content item into the cache thereby to assess how long items are held in the cache before they are expired and deleting the mapping relating to that content item when that content item is judged to have expired in the cache.
4. A method according to claim 1, wherein the step of intercepting traffic is carried out by intercepting traffic flowing out of an original content source.
5. A method according to claim 4, wherein the method further comprises receiving an advertising message which advertises a mapping generated elsewhere on the network and which is related to content items stored in the original content

source, and augmenting the content index using information contained in the advertising message.

- 5 7. A method according to claim 1 wherein the step of intercepting traffic is carried out by intercepting content requests issued by a cache, and wherein the method further comprises advertising the content identities for which mappings are stored in the content index by sending advertising messages to a predetermined location in the network.
- 10 6. A method of retrieving content in an IP-based network comprising the steps of:-

 - 15 (a) intercepting a content request containing information related to the identity of a content item and a specified source location for the content item,
 - (b) sending an instance request to a content index associated with the specified source location, the instance request including the identity of the requested content,
 - (c) receiving an instance mapping from the content index which contains a list of instances and associated locations for the requested content,
 - 20 (d) selecting the best instance of the content from the list,
 - (e) obtaining the requested content from the location associated with the best instance of the requested content, and
 - (f) returning the requested content to the requester of the content.
- 25 7. A proxy for an IP-based network comprising:-

 - (a) a data input operable to receive data from the network,
 - (b) a data output operable to send data to the network,
 - (c) an identity extractor operable to analyse data received at the data input and to extract content identity information from the data,
 - 30 (d) a location extractor operable to analyse data received at the data input and to extract location information from the data,
 - (e) a mapping generator operable to generate a mapping from a content item identified by identity information provided by the identity extractor,

- to at least one destination location identified by associated destination location information provided by the location extractor, and
- (f) a content index database operable to store a mapping provided by the mapping generator and which is operable to provide an instance mapping containing a list of destination locations in response to an instance request containing a content identity.

8. A proxy for an IP-based network comprising:-

- (a) a data input operable to receive data from the network;
- (b) a data output operable to send data to the network;
- (c) a location requester operable to identify a request for a content item in data received at the data input and to send an instance request to a content index associated with the source location of the content item specified in the content request, the instance request including the identity of the requested content, and
- (d) a content returner operable to receive an instance mapping from the content index which contains a list of instances and associated locations for the requested content, to select the best instance of the content from the list, to obtain the requested content from the location associated with the best instance of the requested content, and to return the requested content to the requester of the content.

9. An advertising message for transmission over an IP-based network, the message being arranged to advertise a replica content item and a location for that item.

10. An instance request for transmission over an IP-based network, the request including a destination address and the identity of a requested content item, the destination address being a different address to the address of the source location of the content item.

11. An instance mapping for transmission over an IP-based network, the instance mapping containing a list of instances and associated locations for a predetermined content item.

12. A computer program which when executed indexes content in an IP-based network by:-

- 5 (a) intercepting traffic flowing in the network,
 (b) extracting content identity information and associated destination location information from the traffic flow,
 (c) generating a mapping from a content item identified by the extracted identity information to at least one destination location identified by the associated destination location information, and
 10 (d) storing the mapping in a content index database which is operable to provide an instance mapping containing a list of destination locations in response to an instance request containing a content identity.

13. A computer program which when executed retrieves content in an IP-based network by:-

- 15 (a) intercepting a content request containing information related to the identity of a content item and a specified source location for the content item,
 (b) sending an instance request to a content index associated with the specified source location, the instance request including the identity of the requested content,
 20 (c) receiving an instance mapping containing a list of instances and associated locations for the requested content,
 (d) selecting the best instance of the content from the list,
 25 (e) obtaining the requested content from the location associated with the best instance of the requested content, and
 (f) returning the requested content to the requester of the content.